

Unmanned Aircraft Guidance for Penetration of Pretornadic Storms

Jack Elston* and Eric Frew†
University of Colorado, Boulder, Colorado 80309

DOI: 10.2514/1.45195

This paper investigates unmanned aircraft guidance for a severe storm penetrator to assess the feasibility of such a mission and to optimize ingress in terms of flight time and exposure to precipitation. Understanding guidance-layer behavior and deriving control strategies will provide tools for developing mission-level concepts of operation and making vehicle design tradeoffs. A backward propagating wave front algorithm based on ordered upwind methods is developed for guidance-layer planning of severe storm penetration. Target areas for investigation by the aircraft are identified within simulated storm wind, hail, and precipitation data. The simulated data are used to determine the feasibility of storm penetration to the points of interest, along with calculating airframe exposure to precipitation. Ingress planning is also considered for a dynamic wind field using a receding horizon control approach to adapt the flight path in response to environmental changes.

I. Introduction

ACCORDING to the National Weather Service, 438 fatalities, close to 3000 injuries, and \$11.2 billion in damages were incurred due to severe weather in the United States in 2003 [1]. These losses could be dramatically reduced with effective advanced prediction and warning systems. Tornadoes are especially violent members of the severe storm family and thus, the study of tornado formation and evolution is a public safety necessity. There was an immediate improvement in tornado warning capabilities with the introduction of Doppler radar [2]. However, primitive sensing methods such as human storm chasers and spotters remain the most vital part of the public safety system. The inability to determine the volumetric thermodynamic state of the atmosphere between the ground and the base of the mesocyclone remains a major barrier towards a deeper understanding of tornado genesis. The limitations of remote sensing are evident. One cannot remotely sense the thermodynamic field; these data can only be obtained with in situ sensing.

Research into tornadogenesis will not progress significantly until there are measurements of the thermodynamic and microphysical properties aloft in the vitally important rear-flank region of supercell storms. A consensus of research in the last 25 year makes it clear that a small downdraft of a few kilometers width, known as the rear-flank downdraft, plays a causative role in tornado formation [1,2]. But recent studies have produced a quandary: surface observations from instrumented vehicles beneath this downdraft indicate that it typically arrives at the ground relatively warm and potentially buoyant compared with typical thunderstorm downdrafts, whereas studies of the flow in and around this downdraft suggest that it is negatively buoyant aloft. It is surmised that this negative buoyancy, if present in sufficient quantities upstream of the location of potential tornado formation, causes the rotation that is eventually reoriented and concentrated into a tornado [1,2]. To understand tornado formation scientists must obtain in situ observations of thermodynamics and the types, amounts, and concentration of hydro-

meteors that, through phase changes, influence the thermodynamics. Unfortunately, in situ observations are not possible in and around this downdraft using current technology. Balloons cannot ascend through strong downdrafts and these flows are too dangerous for penetration using manned aircraft, as evidenced by the inadvertent penetration of a rear-flank downdraft [2].

Field observations of supercells and tornado formation occur every spring in the central U.S. Currently these deployments involve two or more Doppler radar stations. The VORTEX-2 experiment will expand this to five or more mobile Doppler radars with the addition of unmanned aircraft systems (UAS) [3]. To develop vehicle specifications and concepts of operation for a VORTEX-2 deployment, the performance of UAS in severe pretornadic storms must be studied. Initial operational constraints point to the use of a small unmanned aircraft (UA) to perform such a task. The UA must be easily portable in a road-worthy vehicle, and must be able to be deployed quickly. This mandates that the UA must be able to be transported fully assembled. Furthermore, cost constraints along with considerations for the potential for vehicle loss also limit its size. Precedence for the use of small UA in the study of atmospheric convection and severe storms has been established by the Aerosonde UA [4–6] using waypoint commands issued by the operator. That UA has been used in several field campaigns to study tropical storms and hurricanes [4–6]. However, to date penetration of a pretornadic storm, which evolves over quicker and smaller scales than tropical storms, has not been reported by any UAS.

This paper investigates the guidance-layer behavior of a severe storm penetrator to assess the feasibility of such a mission and to reduce ingress in terms of flight time and exposure to precipitation. The understanding of guidance-layer behavior and control strategies gleaned from this work is shown to provide tools for developing mission-level concepts of operation and for making vehicle design tradeoffs. For example, aircraft speed generally increases with aircraft mass, so there is a high-level tradeoff between the size of the initial region from which an aircraft can be deployed (due to its ability to fight against strong winds) and the kinetic energy contained in the vehicle (which relates to system safety).

II. Planning Methods

Aircraft guidance in the presence of strong winds presents several challenges, particularly when maximum aircraft speeds are less than the peak wind speeds.[‡] In these cases the aircraft are unable to travel directly upstream in some locations, and so large regions of the

Presented as Paper 6513 at the AIAA Guidance, Navigation and Control Conference and Exhibit, Honolulu, HA, 18–21 August 2008; received 30 April 2009; revision received 20 July 2009; accepted for publication 27 July 2009. Copyright © 2009 by Jack Elston. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/10 and \$10.00 in correspondence with the CCC.

*Graduate Research Assistant, Department of Aerospace Engineering Sciences. Student Member AIAA.

†Assistant Professor, Department of Aerospace Engineering Sciences. Member AIAA.

[‡]The term strong wind field is used to denote a field where the peak wind value is greater than the maximum airspeed of the aircraft

environment may be inaccessible depending on the initial position of the UA. Furthermore, as these wind fields vary over time, the reachable regions of the environment will change. Thus, ingress planning must be able to account for the time-varying field in the configuration space and adapt quickly in response to changes. Because the environment is dynamic and trajectories are likely to change, near-optimal algorithms, such as receding horizon control, that generate plans quickly are preferred to optimal ones that may take significant time.

The guidance-layer planning problems considered in this paper are characterized by the presence of background winds that move the aircraft even when the control inputs are zero, making them examples of control affine systems with drift [7]. In general, techniques such as Lie algebra can be extended to determine controllability [7] of systems with drift. However, these techniques do not address general optimal control for systems with drift. Instead, motion planning algorithms are often used based on specific problem formulations. Recent approaches using wave front expansion via level set [8] and fast marching methods [9] have been applied to linear and control affine systems with drift, especially in the context of autonomous underwater vehicles operating in strong current fields [10].

The wave front expansion methods are well suited to the planning problems considered here. These methods generate solutions quickly by interpolating over discrete grid representations of the environment with approximation error decreasing as the mesh spacing is decreased. In general these methods are not applied to aircraft because wind data are difficult to obtain. However, the Doppler radar systems [3] deployed to study severe storms provide this information at discrete points in the regions of interest and can therefore be incorporated easily by the planning methods. In fact, wind data are only provided at discrete points so planning methods must interpolate between grid points. Aircraft kinematic constraints and environmental constraints due to terrain can also be included in these methods [10].

Wave front expansion algorithms for path planning and optimal control use breadth-first search over a discretized space to approximate the propagation of continuous wave fronts corresponding to solutions of Hamilton–Jacobi equations. A goal region (or point) is indicated and assigned a small value. Each grid point adjacent to the goal is then updated based upon the value of other adjacent points and an interpolated penalty function such as Manhattan or Euclidean distance [11,12]. Once the entire space has been evaluated, a path can be generated from any location to the goal by further interpolating between grid points and using a gradient descent algorithm. Extensions to the basic propagation steps can account for obstacles and other constraints in the configuration space [13]. These types of planning methods require that the entire space be known a priori to generate a path, but only need be called once in a static environment to calculate the optimal trajectory from every point in the environment.

Most wave front expansion methods for path planning through current fields use a forward propagating wave front planner [10,14]. This is well suited for applications that use small amounts of data and are constantly obtaining only local knowledge of the current field with some distant global objective. When considering the problem of storm ingress with a UA, Doppler radar coverage can give a low-resolution global wind field. This enables the generation of a globally optimal ingress path using forward wave front propagation. Unfortunately, real world application of the algorithm will result in differences between the actual and expected trajectories, and will require recalculation of the cost map. Therefore the approach used in this paper uses backward wave front propagation to calculate time-to-go contours from the goal region once per planning cycle (i.e., between sensor updates) to eliminate the need for recomputation at each control sample. The resulting contour map provides information on the feasibility of storm penetration from given locations and yields the ingress trajectory via gradient descent.

The method of wave front expansion also strongly affects the ability to create feasible paths in the presence of strong winds. Some methods for cost propagation limit directions of travel due to the discretized nature of the wind data [10] and can fail to find feasible

paths in strong wind fields even where one exists. A method such as a sliding wave front expansion [14] uses interpretation to provide feasible paths where other algorithms fail, but also only produces one feasible path. Should the path of the vehicle deviate from the optimal path, a new forward propagation must be performed. Ordered upwind methods [15] address both feasible path limitations due to coarse granularity of sensor data and uses backward propagation. It accommodates the previously mentioned limitations by continuously refining cost values on the fringe of the wave front.

Several fundamental issues remain to be addressed before existing methods can be applied to aircraft guidance. First, the presence of a strong wind field can restrict the motion of the aircraft. There will be regions in the environment where the aircraft cannot head against the wind. Second, the wind field leads to anisotropic drift terms that complicate the wave front expansion. Fast marching methods for wave front expansion are designed for isotropic wave fronts that expand with equal speed in all directions. When these methods can be applied they are very efficient and computationally fast. Unfortunately they cannot be applied directly here. Third, existing methods have only been applied to static fields [10,14]. Time-varying fields complicate the wave front expansion because the goal region can take on a distorted shape in the configuration space. These issues are addressed here by the proper development of a penalty cost function, ordered upwind wave front expansion, and receding horizon planning, respectively.

This paper presents a backward propagating wave front algorithm based on ordered upwind methods [15] for guidance-layer planning of severe storm penetration by unmanned aircraft. Feasibility analysis is performed on a static wind field to assess baseline performance expectations. Ingress planning is then considered for a dynamic wind field using a receding horizon control approach to adapt the flight path in response to environmental changes.

II. System Models

A. Unmanned Aircraft Dynamics

The UA model used here assumes the presence of a low-level flight control system capable of altitude-hold, speed-hold, and turn rate-command functions. For basic analysis we assume the flight control system presents to the guidance layer the standard kinematic model

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} V_a \cos \psi + W_x(x, y, z, t) \\ V_a \sin \psi + W_y(x, y, z, t) \\ W_z(x, y, z, t) \\ u \end{bmatrix} \quad |u| \leq \omega_{\max} \quad (1)$$

where V_a is the constant aircraft airspeed, u is the turn rate command, ω_{\max} is the maximum allowable turn rate, and $[W_x, W_y, W_z]^T$ are the position- and time-dependent wind components derived from storm data given at discrete grid points. For UA locations between grid points, cubic interpolation (using MATLAB's interpolation method) is used to determine the wind speed. Although vertical motion is considered in this work, only planar (i.e., in the x - y plane) motion is controlled. Therefore, to simplify notation let the aircraft planar velocity be $\mathbf{v} = [\dot{x}, \dot{y}]^T$, let the planar wind velocity be $\mathbf{v}_w = [W_x, W_y]^T$ with $V_w = \|\mathbf{v}_w\|$, and let the aircraft velocity relative to the wind be $\mathbf{v}_a = [V_{a,x}, V_{a,y}]^T = [V_a \cos \psi, V_a \sin \psi]^T$. Note that $V_a = \|\mathbf{v}_a\|$ and $\mathbf{v} = \mathbf{v}_a + \mathbf{v}_w$.

B. Storm Data

Unmanned aircraft performance is investigated through flight in several characteristic severe storms. The storm data sets used here were provided by Jerry Straka of the School of Meteorology at the University of Oklahoma [16–18]. The first set is planar two-dimensional data for a single instant in time for an altitude of 500 m. Storm data are given on a 210 by 210 grid with a resolution of 500 meters between grid points in a reference frame that moves with the storm (with velocity [18, 4] m/s). Data include rain, hail, and three-dimensional wind velocity components (relative to the storm frame of reference). The storm simulation is aligned such that the

positive y axis corresponds to north. Figure 1 shows contour plots of the storm data including rain, hail, planar (x - y) wind speed, and the z -component of wind vector. Figure 2 shows a close-up of the main portion of the rain contour overlaid with planar wind velocity vectors.

The second set was generated using different initial conditions on a three-dimensional grid at 1 min time intervals. The spatial representation is given on a 100 by 100 by 25 grid where the grid points have separations of 500, 500, and 100 m between points (in the x , y , and z directions, respectively). Data are given every minute for 16 min, representing the full evolution of the storm.

The rear-flank downdraft, which is the goal region for pretornadic storm penetration (see Fig. 2), is characterized by several traits, with some observable in the simulation data. First, as the name implies, the wind velocity in the region has a strong downward component (Fig. 1). Second, the region sits behind a gust front that is observable in the wind data (Fig. 2). Third, the region contains a large amount of rain (also observable in Fig. 2). Finally, the rear-flank downdraft sits behind a Cumulonimbus cloud base that is not visible in the data, but is visible in the field (and marked on Fig. 2). These traits of the rear-flank downdraft allow both the operator and the ingress planning system to identify and track the moving goal region during the penetration mission.

III. Path Planning Using Wave Front Expansion Methods

Guidance-layer path planning algorithms for storm penetration are developed based on ordered upwind methods for wave front propagation. Ordered upwind methods (OUM) are a class of noniterative algorithms that solve continuous wave front propagation problems on a discretized grid representation of the configuration space. These priority-based schemes maintain an ordering of points and systematically compute the solution by relying on known, previously computed information that lies upstream, that is, information that would be visited by the expanding wave front first [15]. These methods provide an approximation to the wave front propagation by interpolating between discrete grid points, with the approximation known to converge to the optimal solution in the limit as grid spacing decreases [15]. In the context of path planning and optimal control,

OUM is used to approximate the solution of the Hamilton–Jacobi equations that describe optimal aircraft trajectories. The ordered upwind methods use partial information about the characteristic directions of the wave front propagation to decouple the nonlinear systems described by the Hamilton–Jacobi equations, producing one-pass algorithms of greatly reduced computational labor compared with other numerical techniques [15].

The planning problem of interest here is to minimize the objective

$$J = \int_{\mathbf{x}_0}^{\mathbf{x}_f} g(\mathbf{x}) d\mathbf{x} \quad (2)$$

where \mathbf{x}_0 and \mathbf{x}_f are the initial and final (goal) positions, respectively, of the aircraft and $g(\mathbf{x})$ is some cost function. Taking $g(\mathbf{x}) = \dot{\mathbf{x}}^{-1}(\mathbf{x})$ makes the objective the duration of the path and leads to the time-optimal trajectory generation problem. When the aircraft path is

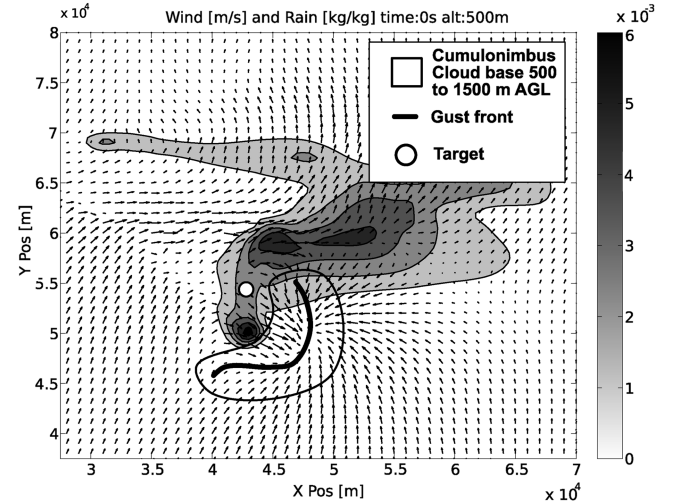


Fig. 2 Contours of rain with planar wind vectors around region of interest. Sensing target, gust front, and cloud base have been indicated by hand.

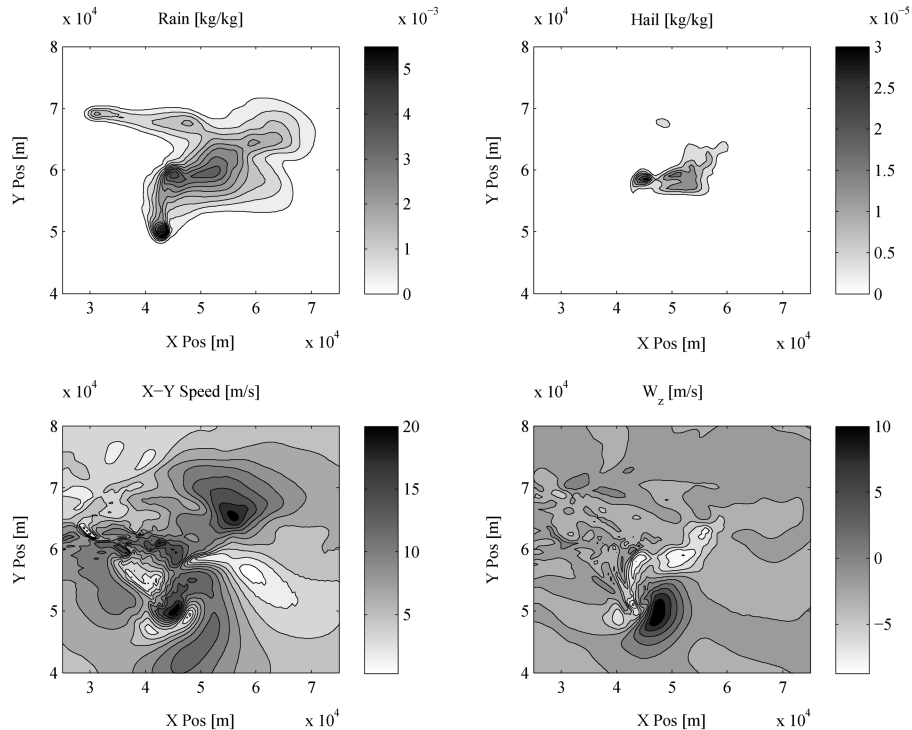


Fig. 1 Contours of rain, hail, 2-D speed, and z component of wind speed for a simulated storm environment. Rain and hail values are given as the ratio of the mass of the precipitation in a particular mass of air, [kg/kg].

denoted by a discrete set P_a of ordered points \mathbf{x}_i then the objective can be rewritten as

$$J = \sum_{i=1}^{|P_a|-1} c(\mathbf{x}_i, \mathbf{x}_{i+1}) \quad (3)$$

where

$$c(\mathbf{x}_i, \mathbf{x}_{i+1}) = \int_{\mathbf{x}_i}^{\mathbf{x}_{i+1}} g(\mathbf{x}) d\mathbf{x} \quad (4)$$

is the cost of moving from one position to the next.

Minimizing Eq. (3) subject to the kinematic model Eq. (1) is difficult, and so we plan paths using a less restrictive first-order model and then use a feedback control law to follow the path. This is a rough approximation, but one that the authors do not feel will compromise the effectiveness of the planner, especially given the large grid spacing of the sampled environmental data with respect to the turn radius of the vehicle. It is left to further work to analyze the effect of this simplification. Using the first-order model the aircraft motion is given by

$$\dot{\mathbf{x}} = v(\mathbf{x}, \mathbf{a}) \mathbf{a} \quad (5)$$

where \mathbf{x} is the aircraft position, $v(\mathbf{x}, \mathbf{a})$ is the aircraft speed, and \mathbf{a} is the unit vector that denotes the direction of the velocity. The speed is a function of both the aircraft position and direction. It will be shown in the next section that even when the aircraft's airspeed is constant $v(\mathbf{x}, \mathbf{a})$ will change as a function of position and direction based on the background wind velocity. Finally, the control input to the system is the direction of the velocity vector $\mathbf{u} = \mathbf{a}$, and so the planning problem of interest is to solve for

$$\mathbf{u}^*(\cdot) = \arg \min_{\mathbf{u}(\cdot)} \int_{\mathbf{x}_0}^{\mathbf{x}_f} g(\mathbf{x}) d\mathbf{x} \quad (6)$$

subject to Eq. (5).

A. Cost Function

Using a cost function that ensures that all resulting paths are feasible (i.e., that the aircraft can move against the wind in a certain direction) is pertinent to the success of the presented mission. Previous cost functions, although acceptable given there are no strong winds ($V_a \gg V_w$), can yield infeasible paths in strong winds [10]. Given the design constraints imposed on the unmanned aircraft, it is assumed that the environment encountered by a UA when approaching a pretornadic storm will contain winds that will likely be larger than the maximum air speed of the vehicle.

To ensure that all generated paths are feasible, travel time for the vehicle to reach a target point needs to be derived correctly [14]. To generate the cost function it is assumed that the aircraft can move in any direction with speed V_a and that a path following controller can be implemented such that the aircraft described by Eq. (1) can stay on the resulting path. Given a goal point a distance \mathbf{d} away, aircraft velocity \mathbf{v}_a , and wind field velocity \mathbf{v}_w , the travel time τ is related to these variables by $\mathbf{d} = (\mathbf{v}_w + \mathbf{v}_a) \cdot \tau$. Projecting this on the x and y axis yields

$$d_x = (V_{a,x} + W_x) \cdot \tau \quad d_y = (V_{a,y} + W_y) \cdot \tau \quad (7)$$

Aircraft speed components $V_{a,x}$ and $V_{a,y}$ can be eliminated by substituting $V_a^2 = V_{a,x}^2 + V_{a,y}^2$ into Eq. (7) to give

$$(d_x - W_x \cdot \tau)^2 + (d_y - W_y \cdot \tau)^2 = V_a^2 \cdot \tau^2 \quad (8)$$

Solving this quadratic equation for travel time yields

$$\tau = \frac{-(W_x \cdot d_x - W_y \cdot d_y) + \sqrt{\Delta}}{V_a^2 - V_w^2} = \frac{\sqrt{\Delta} - \langle \mathbf{d} \cdot \mathbf{v}_w \rangle}{V_a^2 - V_w^2} \quad (9)$$

where $\Delta = V_a^2 \cdot (d_x^2 + d_y^2) - (W_x \cdot d_y - W_y \cdot d_x)^2$. Considering the case where $V_a = V_w$, the equation can be simplified to:

$$\tau = \frac{d^2}{2(\mathbf{d} \cdot \mathbf{v}_w)} \quad (10)$$

The time τ is only defined when $\Delta \geq 0$. The union of the set where τ is undefined, and when $\tau < 0$ identifies infeasible regions of travel.

When solving the time-optimal planning problem, the cost for moving from position \mathbf{x}_i to \mathbf{x}_j is defined as

$$c(\mathbf{x}_i, \mathbf{x}_j) = c_{ij} = \tau \quad (11)$$

with $\mathbf{d} = \mathbf{x}_j - \mathbf{x}_i$. If it is assumed that exposure to precipitation has an upper bound for safe aircraft flight then precipitation can be accounted for by treating rain and hail as obstacles. In this case the cost function is modified

$$c_{ij} = \begin{cases} \tau: \text{rain}(x, y, z, t) \leq \text{rain}_{\max} & \text{and} & \text{hail}(x, y, z, t) \leq \text{hail}_{\max} \\ \infty: \text{rain}(x, y, z, t) > \text{rain}_{\max} & \text{or} & \text{hail}(x, y, z, t) > \text{hail}_{\max} \end{cases} \quad (12)$$

When the cumulative (additive) effect of precipitation needs to be included in the optimization, a second term is added to the cost function

$$c_{ij} = \begin{cases} \tau + \beta_p(\mathbf{x}_i) \cdot d: \text{rain}(\mathbf{x}_i) \leq \text{rain}_{\max} & \text{and} & \text{hail}(\mathbf{x}_i) \leq \text{hail}_{\max} \\ \infty: \text{rain}(\mathbf{x}_i) > \text{rain}_{\max} & \text{or} & \text{hail}(\mathbf{x}_i) > \text{hail}_{\max} \end{cases} \quad (13)$$

with

$$\beta_p(\mathbf{x}) = k_{\text{rain}} \cdot \text{rain}(\mathbf{x}) + k_{\text{hail}} \cdot \text{hail}(\mathbf{x}) \quad (14)$$

where $\text{rain}(\mathbf{x})$ and $\text{hail}(\mathbf{x})$ are the precipitation values at \mathbf{x} and k_{rain} and k_{hail} are gains that can be used to tune the optimization. Using the cost function in Eq. (13) is equivalent to solving Eq. (6) with $g(\mathbf{x}) = \dot{\mathbf{x}}^{-1}(\mathbf{x}) + \beta_p(\mathbf{x})$.

B. Wave Front Expansion

Wave front expansion is performed using ordered upwind methods to propagate cost-to-go values over the discretized configuration space of the unmanned aircraft. Unlike standard fast marching expansion methods, OUM accounts for the anisotropic nature of the wind field during wave front propagation. Ordered upwind methods are implemented via finite difference update rules that are proven to converge in the limit as grid spacing decreases to the viscosity solution of the Hamilton–Jacobi partial differential equation corresponding to the optimal control problem, defined here based on the cost function in Eq. (13), [15].

The optimal cost-to-go map is generated by dividing the considered space into three sets of points, far, considered, and accepted (the Nomenclature and description presented here closely follow that presented in [15]). The accepted front is defined to be the set of accepted points adjacent to at least one member of the set of considered points. Let AF be the set of line segments $\mathbf{x}_j \mathbf{x}_k$ where \mathbf{x}_j and \mathbf{x}_k are grid points on the accepted front such that there is a considered grid point \mathbf{x}_i adjacent to both points. For each considered grid point \mathbf{x}_i the set of points on the accepted front near \mathbf{x}_i is defined

$$NF(\mathbf{x}_i) = \left\{ (\mathbf{x}_j, \mathbf{x}_k) \in AF \mid \exists \mathbf{x} \text{ on } (\mathbf{x}_j, \mathbf{x}_k) \text{ s.t. } \|\mathbf{x} - \mathbf{x}_i\| \leq h \frac{F_2}{F_1} \right\} \quad (15)$$

where h is the grid size and F_1 and F_2 are the upper and lower bounds on the speed of the AF , respectively, and make up the anisotropy ratio $\Gamma = F_2/F_1$. Let $U(\mathbf{x}_i)$ represent the optimal cost to go from point \mathbf{x}_i , let $V(\mathbf{x}_i)$ be a candidate value for the optimal cost to go from point \mathbf{x}_i , and let $V_{\mathbf{x}_j, \mathbf{x}_k}(\mathbf{x}_i)$ be a consistent upwinding approximation when the characteristic \mathbf{x}_i lies in the simplex $\mathbf{x}_i \mathbf{x}_j \mathbf{x}_k$, which we take here to be

$$V_{\mathbf{x}_j, \mathbf{x}_k}(\mathbf{x}_i) = \min_{\zeta \in [0,1]} \{ \zeta [c_{ij} + U(\mathbf{x}_j)] + (1 - \zeta) [c_{ik} + U(\mathbf{x}_k)] \} \quad (16)$$

where c_{ij} is the value of the cost function when moving from \mathbf{x}_j to \mathbf{x}_i , for example, Eq. (13). Finally, let Ω be the domain representing the environment and $\delta\Omega$ be the goal region in the environment.

The following is the method for generating a cost map using the ordered upwind method:

- 1) Start with all grid points in far ($V(\mathbf{x}_i) = \infty$).
- 2) Move the boundary grid points ($\mathbf{x}_i \in \partial\Omega$) to accepted [$U(\mathbf{x}_i) = 0.0$].
- 3) Move all the grid points \mathbf{x}_i adjacent to the boundary into considered and evaluate the tentative value of $V(\mathbf{x}_i) = \min_{(\mathbf{x}_j, \mathbf{x}_k) \in NF(\mathbf{x}_i)} V_{\mathbf{x}_j, \mathbf{x}_k}(\mathbf{x}_i)$.
- 4) Find the grid point \mathbf{x}_r with the smallest value of $V(\mathbf{x}_i)$ among all the considered.
- 5) Move \mathbf{x}_r to accepted and update the accepted front [$U(\mathbf{x}_r) = V(\mathbf{x}_r)$].
- 6) Move the far grid points adjacent to \mathbf{x}_r into considered.
- 7) Recompute the value for all the considered \mathbf{x}_i within the distance $h \frac{F_2}{F_1}$ from \mathbf{x}_r . If the new computed value is less than the previous tentative value for \mathbf{x}_i then update $V(\mathbf{x}_i)$.
8. If considered is not empty then go to 4.

Figure 3a shows a test case that was shown to fail in an isotropic wave front propagation given certain vehicle speeds due to the fact that the corridor of feasibility does not intersect any grid points [14]. Figure 3b shows the optimal cost-to-go map generated from a wave front using the ordered upwind method. It is apparent from the figure that a feasible solution can be found by descending the gradient of this map.

IV. Application

The ordered upwind methods described in the previous section are applied to the simulated storm data. The results demonstrate how OUM 1) provide a valuable tool for determining the feasibility of UA ingress, 2) enable feedback control for aircraft guidance, and 3) can be used for path planning during ingress.

A. Feasibility, Optimal Ingress Time, and Aircraft Design

Feasibility and ingress time can be evaluated by using Eq. (9) as the penalty function in the ordered upwind expansion process. The result of the OUM using this cost gives an approximation of the optimal time to go from every point in the environment. For locations in the environment that are not grid points, the time to go is determined by interpolation. Regions from which the aircraft cannot possibly reach the goal location will be denoted by an undefined (or very large) value in the map.

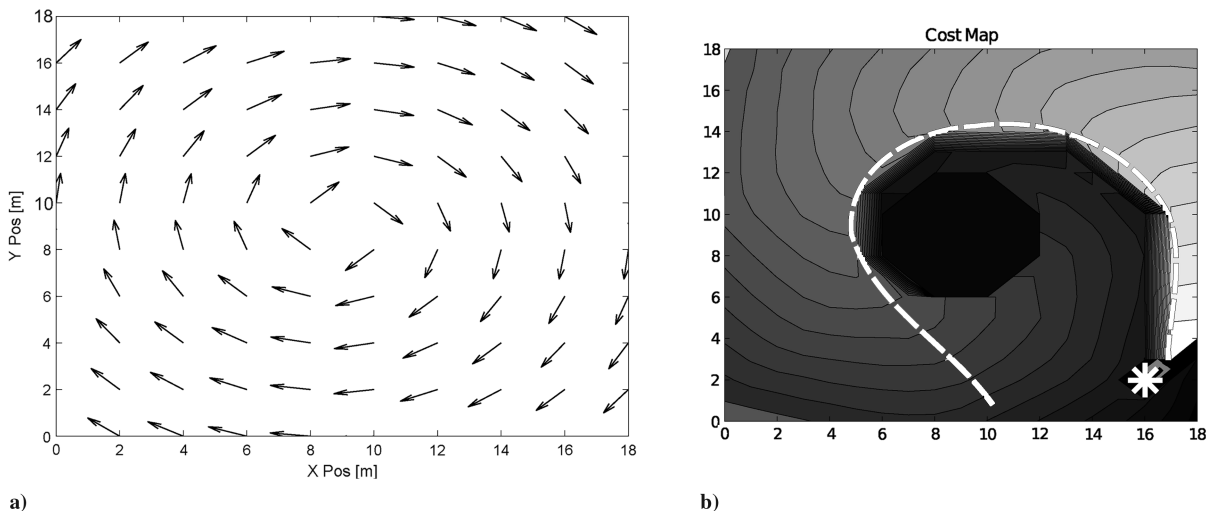


Fig. 3 a) Test wind field, constant speed of 30 m/s, and b) cost map obtained by using an ordered upwind method for backward wave front expansion with $V_a = 20$ m/s. An example path generated using this cost map to a goal location indicated by the asterisk is shown as a dashed line.

Figure 4 contains examples of the time-to-go maps generated using simulated storm data (Fig. 4a) with the goal location at $\mathbf{x}_{\text{goal}} = [4.25, 5.4]^T \cdot 10^4$ m and aircraft with constant air speeds of $V_a = \{7.5, 10.0, 20.0\}$ m/s (Figs. 4b and 4c). For a slow moving aircraft (Fig. 4a), the wind field in the storm separates the environment into two distinct regions: the southern region where the aircraft is able to reach the goal location quickly (within 10 min or less) and the northern region where ingress time is significantly larger (30 min or greater). Note, the simulation extends beyond the borders of Fig. 4a, which just shows results near the goal region, so that trajectories from the northern region can in fact flow down into the southern region. The next plot (Fig. 4b) shows that increasing the airspeed slightly to $V_a = 10.0$ m/s changes the time-to-go map significantly. Ingress paths from the southern region are opened up along corridors on the eastern and western edges. The map further shows that the western corridor is generally better because the aircraft flies with the wind field there. The center of the environment is still marked by a boundary that cannot be crossed. Finally, for air speed of $V_a = 20.0$ m/s, Fig. 4c shows that the aircraft can fight the wind from any location to reach the goal without having to go around the center of the environment. The time to go is still varied across the environment, but all paths are reasonably direct.

The time-to-go maps can be used in several different ways. As a design tool, the maps can be used to evaluate the tradeoff between aircraft speed (which usually relates to mass) and mission-level performance objectives like ingress time and deployment feasibility. Figure 4 shows wide variation in the general shape of the time-to-go map, implying very different performance results for different aircraft airspeeds. A salient example is the map with $V_a = 7.5$ m/s, which shows that the entire northern half of the environment has trajectories at least 3 times slower than the southern region. Thus, if mission-level constraints are going to force the aircraft to deploy north of the storm, the airspeed should be increased.

A second use of the cost-to-go maps is the selection of deployment locations. The figures show that given vehicles with slower speeds, deployment location becomes more important. In some cases small movements in deployment location (less than 5 km) can significantly change the time for the UA to reach the goal location. Consider the map for $V_a = 10.0$ m/s (Fig. 4c), which shows two corridors along the edges. Deployment from $[5.5, 6.2]^T \cdot 10^4$ m leads to a path that has to travel around the center and into the western corridor, taking approximately 15 min to reach the goal. Moving the deployment to the location $[5.5, 6.3]^T \cdot 10^4$ m puts the aircraft on the other side of the central obstruction and lets the aircraft head straight toward the goal region, reaching it in close to 8 min. From an operational perspective, the time-to-go map implies that mission-level optimization should be performed that combines time required to reach a deployment location as well as airborne time to go.

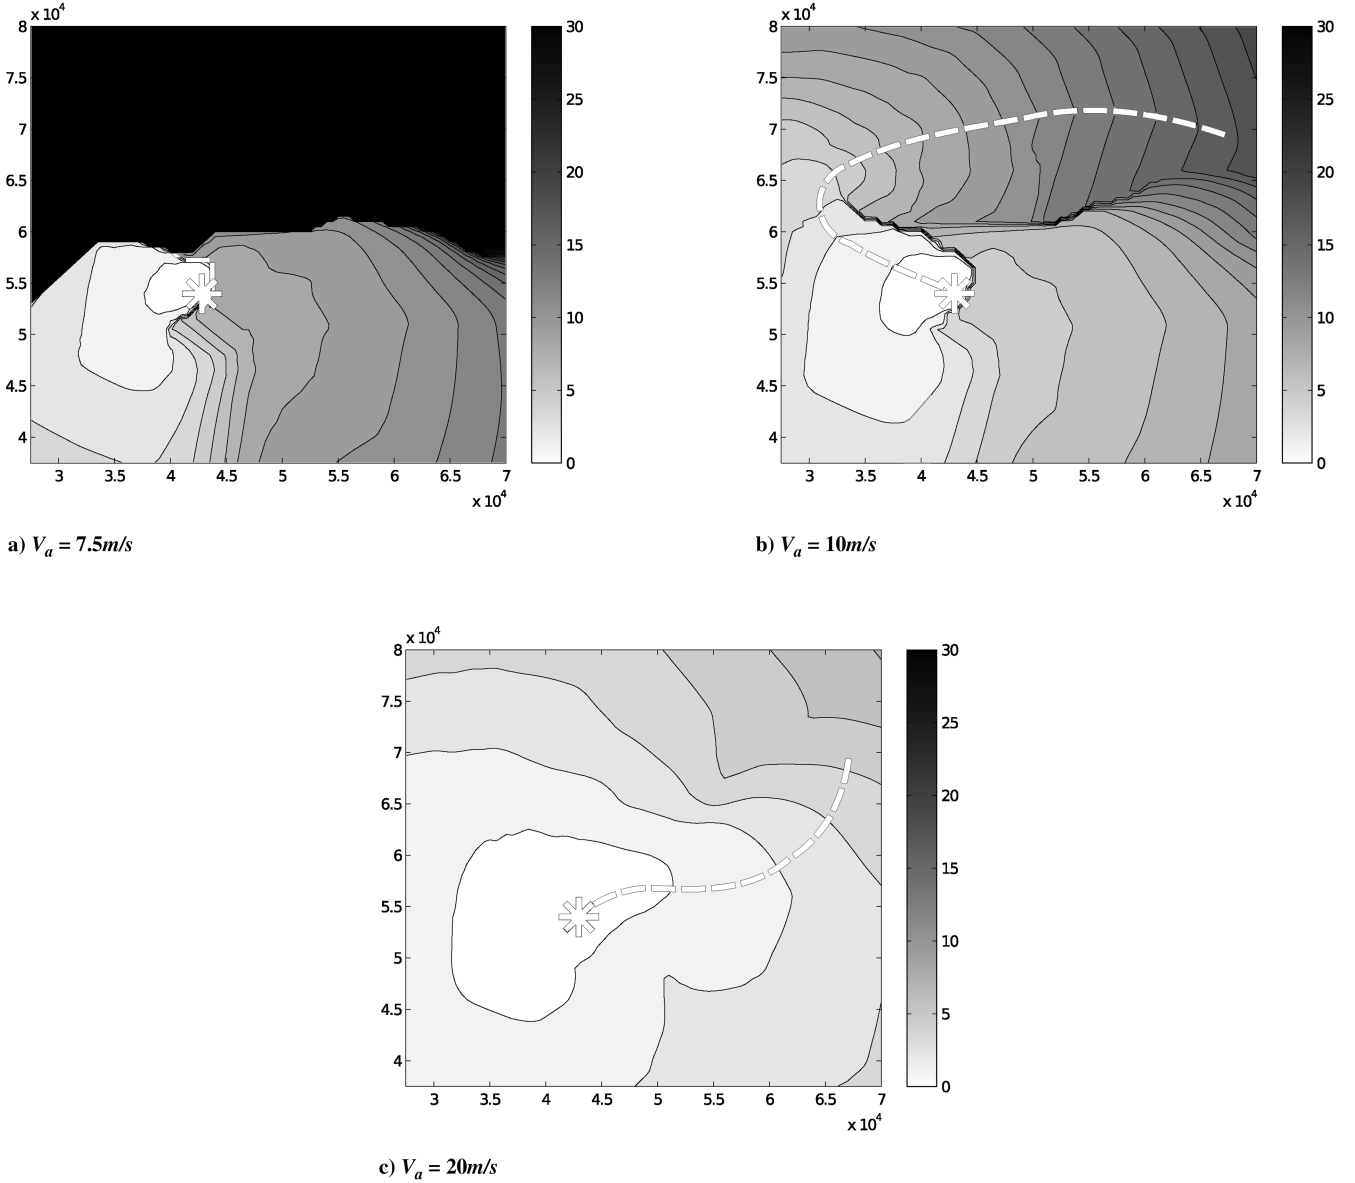


Fig. 4 Cost maps obtained using an ordered upwind method for backward wave front expansion, target indicated by an asterisk. Shading correlates with time in minutes required to reach goal location. Each plot depicts a different vehicle speed, and an example path from the same starting location is shown on each plot as a dashed line. Plot (a) does not contain a path, as it is impossible to reach the goal location within the 30 min time frame.

The cost-to-go maps can also be used to derive the ingress trajectory. A key trait of OUM is that the propagated wave front is an approximation of the viscosity solution of the Hamilton–Jacobi equations associated with the optimal trajectory design problem [15]. The trajectory is determined by following the gradient of the time-to-go map. Once the map is determined for a given goal and wind field, it can be used at any time to determine a feasible and nearly optimal trajectory from any location. From an operational perspective, this enables instantaneous response once the deployment location is determined. Having the complete map also allows for instantaneous adaptation when disturbances push the aircraft off the optimal trajectory. In fact, the next section shows how the time-to-go map can be viewed as a feedback control law that can be used for adaptive ingress planning.

Another example shows how the OUM can account for precipitation. Figure 5 shows the resulting cost-to-go map derived using Eq. (13) with a rain exposure limit $\text{rain}_{\max} = 4 \cdot 10^{-3} \text{ kg/kg}$ ($\text{hail}_{\max} = \infty$). In this example the rain exposure limit creates two obstacles that are denoted as infeasible regions in Fig. 5. The resulting paths move away from these regions before curving toward the goal. The parameter k_{rain} was heuristically set to weight the relative importance of minimum time and minimum rain exposure for this scenario.

B. Ingress Planning

Ingress planning can be realized using the cost maps generated previously for feasibility studies. By following the gradient of the cost map, the UA will reach the goal location in the minimal amount of time (within the limits of the sampled data and assumptions required by the algorithm). Given the map, the path to the goal point can be defined by the set of points \mathbf{x}_t , $t = 1, \dots, n$ where t is the discrete time, x_0 is the deployment location of the UA, and

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \Delta x \cdot \mathbf{n}_t \quad (17)$$

where $\Delta x = T_s / \|\nabla U(\mathbf{x}_t)\|$ is the step size, $\mathbf{n}_t = \nabla U(\mathbf{x}_t) / \|\nabla U(\mathbf{x}_t)\|$ is the step direction, and T_s is the sample time of the discretization. Although knowing the complete optimal trajectory can be beneficial for higher level reasoning, it is not necessary for aircraft guidance. Instead, the aircraft at \mathbf{x}_t simply uses Eq. (17) directly to determine the local direction of its next step.

In practice the aircraft will not follow the optimal direction but will instead move to

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \Delta x \cdot \mathbf{n}_t + \mathbf{e}_t \quad (18)$$

where \mathbf{e}_t is a disturbance vector due to the fact that the wind field measurements will have errors, the goal location may be moving, and




Fig. 5 Cost map obtained by using an ordered upwind method for backward wave front expansion, target indicated by asterisk. Cost function includes a rain exposure limit $\text{rain}_{\max} = 4 \cdot 10^{-3} \text{ kg/kg}$ ($\text{hail}_{\max} = \infty$). Fill lightness correlates with time in minutes to reach goal location given $V_a = 10 \text{ m/s}$, black indicates a boundary. An example path to the goal location is depicted as a dashed line.

the value function $U(\mathbf{x}_t)$ calculated by the OUM is an approximation. In this case the aircraft simply uses the time-to-go map at the new location and moves based on $\nabla U(\mathbf{x}_{t+1})$.

In an operational deployment Doppler radar systems will provide wind field estimations and storm structure at irregular intervals. Let T_D represent the time between successive wind field measurements and let U^k be the cost-to-go map determined from wind field data collected at time k . In this case it is expected that the wind field changes over time so the value function U^k for $k \leq t < k + T_D$ will accrue errors as the actual wind field evolves. The feedback control structure of the backward propagation solution is used as part of a receding horizon controller for ingress planning. In particular, over the time interval $k \leq t < k + T_D$ the aircraft moves in the direction of $-\nabla U^k(\mathbf{x}_t)$. At time $t = k + T_D$ new measurements of the wind field are collected and the new cost-to-go map U^{k+T_D} is calculated. Thus, the aircraft uses the receding horizon control

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \frac{T_s \cdot \nabla U^k(\mathbf{x}_t)}{\|\nabla U^k(\mathbf{x}_t)\|^2} + \mathbf{e}_t, \quad k \leq t \leq k + T_D \quad (19)$$

Depending on model information available to the guidance layer, an additional term $\hat{\mathbf{e}}_i$ that estimates the motion of the target region can be added to the update rule

$$\mathbf{x}_{i+1} = \mathbf{x}_i - \Delta x \cdot \mathbf{n}_i + \mathbf{e}_i - \hat{\mathbf{e}}_i \quad (20)$$

where the estimate $\hat{\mathbf{e}}_i$ is derived from a simple update law like

$$\hat{\mathbf{e}}_{i+1} = \hat{\mathbf{e}}_i + \gamma(\mathbf{x}_{\text{opt},i+1} - \mathbf{x}_{i+1}) \quad (21)$$

where γ is a weight used to determine the frequency response of the update and $\mathbf{x}_{\text{opt},i+1} = \Delta x \cdot \mathbf{n}_i$ is the location the aircraft would have moved to had the wind field been exact.

Thus far derivation of the ingress planner has assumed that the aircraft can travel instantaneously in the direction of $-\nabla U^k$. Unfortunately the kinematic aircraft is constrained to move in the heading direction only with a bounded turn rate. In this case an additional proportional control law is wrapped around the ingress planner to calculate the aircraft turn rate $u = \psi_{\text{opt}} + k_u(\psi_{\text{opt}} - \psi)$ where $\psi_{\text{opt}} = \angle[-\nabla U(\mathbf{x}_t)] = \angle(\mathbf{n}_t)$ is the angle of the gradient of the cost-to-go map relative to the x axis. Again, because the backward propagation using OUM provides $U^k(\mathbf{x}) \forall \mathbf{x}$, the resulting ingress planner performs like a feedback control law.

Two simulations, one with each data set, were conducted to demonstrate the ingress planner. For the two-dimensional set, the UA

was assumed to maintain the same altitude throughout the flight. Because the storm data are given in a moving reference frame, their location over time with respect to the inertial reference frame was set to move with the constant velocity $[18, 4]^T \text{ m/s}$. A linear interpolation between data points in both space and time is used to generate the actual wind field that affects the dynamics of the vehicle. The onboard planner only receives wind field updates at the rate of once per minute, corresponding to a simulated update from Doppler radar information. From these updates the planner calculates the cost map used for guidance for the next minute.

Figure 6 shows the track of the aircraft for the entire simulation and the final location of the storm rain contours. The points where the onboard radar data were updated are marked as triangles. At each of these points, the path generated from the cost-to-go map is shown as a dashed line, and the corresponding goal location is shown as a circle. This goal location was determined by examining the data and choosing the area of interest by hand. The initial path generated by the ingress planner curves toward the left to move out of adverse winds first and then ride favorable winds to the target. As the target region moves, the planner modifies the path in response. After intercept of the target, the vehicle exhibits an oscillatory behavior about the target location. The cause of this behavior is the path following algorithm. Given that an aircraft requires a minimum forward airspeed, once the target is reached the UA is commanded to orbit at a specified radius. The beginning of an orbit combined with periodic radar updates of the target location causes the vehicle to oscillate.

Figure 7 shows the error between the expected location of the aircraft and the actual position during the simulation. These errors reflect several factors, mainly the difference between the path planned from the $\frac{1}{60} \text{ Hz}$ radar update and the path affected by spatially and temporally interpolated simulation data. The error is also affected by other simplifications and assumptions already specified earlier in this paper. The error value in the graph generally increases until a radar update is received by the controller. This behavior validates the usefulness of the addition of an update law [Eq. (20)], which was not included in this simulation, to the controller.

Simulation of flight through the four-dimensional data was performed similarly to the previous simulation. The full data set was interpolated to determine the actual environment for the UA at a given point in time and space. The path planner only received updates on the wind and precipitation fields once a minute to simulate a radar data feed. The UA altitude was allowed to change due to the wind field,

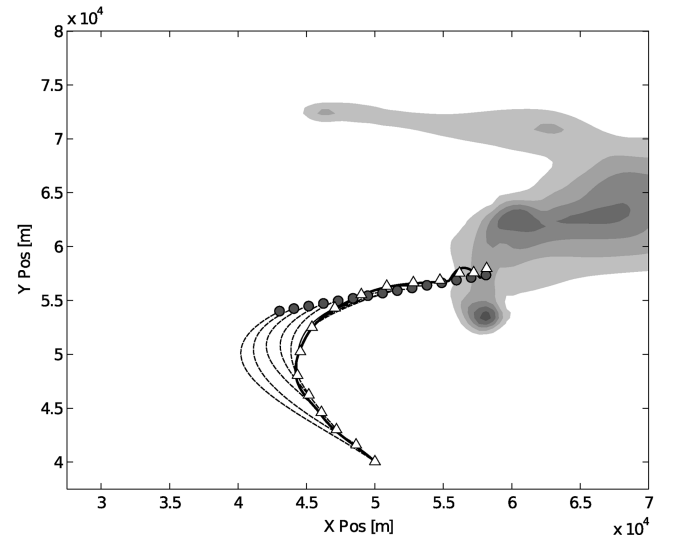


Fig. 6 Rain contours for the first set of storm data at the end of the simulation. The track of the vehicle is shown as the thick black line, with triangles at the points where the radar data were updated. Circles show the corresponding goal points at each update. The dotted lines show the calculated optimal path at the start each iteration.

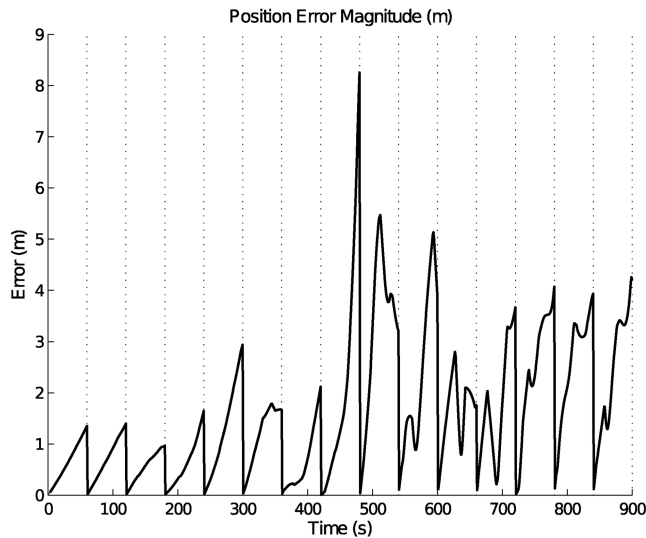


Fig. 7 Error between expected and actual position during the simulation.




Fig. 8 Rain contours for the second set of storm data at the end of the simulation. The track of the vehicle is shown as the thick black line, with triangles at the points where the radar data were updated. Circles show the corresponding goal points at each update. The dotted lines show the calculated optimal path at the start each iteration.

but was not actively controlled. To reduce the large amount of computation required for producing this simulation, only a two-dimensional slice of the storm data nearest the current altitude of the UA was used for creating the cost maps. Because of this simplification, three-dimensional control could not be accommodated by the path planner. It is possible that a simplification of the data would have allowed for three-dimensional control to be considered, but it is left to future work to determine the computational constraints of the algorithm, their effect on the applicability of the algorithm, and the possible solutions.

Figure 8 shows the rain contours for the sixteenth minute of the data set with an overlaid UA track. The initial UA path moves upward before turning left to head into the rear-flank downdraft. This path takes it above the gust front with large downdrafts. As the storm moves downward, the UA begins to encounter more rain and wind than expected. After the second planning epoch, the optimal path shifts to a southern route below the gust front.

V. Conclusions

Ordered upwind methods were used to develop guidance-layer planning algorithms for severe storm penetration by small unmanned aircraft. The ordered upwind methods are a class of noniterative

algorithms that solve continuous wave front propagation problems on a discretized grid with significantly less computational effort than other numerical techniques. In the context of path planning and optimal control, ordered upwind methods approximate the solution of the Hamilton–Jacobi equations that describe optimal aircraft trajectories. They are well suited to the planning problems considered here because they are designed for anisotropic flowfields and because the Doppler radar systems deployed to study severe storms provide wind field information at discrete points in the environment.

A backward propagating wave front algorithm based on ordered upwind methods was created to aid development of mission-level concepts of operation for storm penetration. In particular, the algorithm generates cost-to-go maps of the environment that 1) can be used to determine the feasibility of storm penetration from different locations, 2) can be used as a design tool to understand aircraft speed requirements, 3) show the sensitivity of ingress time to deployment location, 4) yield feasible ingress trajectories that reduce flight time, and 5) form a component of a receding horizon ingress planner. Unlike previous work that uses forward propagating wave front planners in current fields, the global cost-to-go maps generated here using backward propagation are used as feedback control laws to simplify the overall computational of the guidance-layer inputs.

The planning algorithms were applied to simulated severe storm data to highlight their utility. The simulated data emulate measurements that can be taken in the field by Doppler radar systems. Several different examples showed how storm winds create distinct regions in the environment such that vehicle design and deployment location can have significant impact on overall mission performance. The inclusion of precipitation as a constraint and as an additional term in the optimization objective function were shown. Finally, ingress planning in the face of unknown disturbances was also demonstrated.

Future work will address several issues. We are investigating additional modifications to the ordered upwind methodology to account for path constraints from vehicle dynamics and the cumulative effect of the vertical wind component. Further, adaptive mesh generation is being pursued to reduce the overall computational complexity of ingress planning. Finally, estimation from reduced order storm models will be incorporated into the ingress planner.

Acknowledgments

This work was supported by the National Science Foundation under awards ATM-0715941 and ATM-0824160. The authors would like to thank Jerry Straka and Erik Rasmussen for providing the simulated storm data and helping us interpret it. The authors also thank Adam Houston and Brian Argrow for supporting the overall objectives of this effort.

References

- [1] National Weather Service Office of Climate Weather and Water Services, "NWS Weather Fatality, Injury, and Damage Statistics," <http://www.nws.noaa.gov/om/hazstats.shtml>, 2008.
- [2] Argrow, B., Lawrence, D., and Rasmussen, E., "UAV Systems for Sensor Dispersal, Telemetry, and Visualization in Hazardous Environments," *43rd Aerospace Sciences Meeting and Exhibit*, AIAA, Reston, VA, Jan. 10–13 2005.
- [3] "VORTEX-2 Website," <http://www.vortex2.org/>, 2008.
- [4] Hipskind, S., Tyrell, G., Holland, G., and Curry, J., "Use of the Aerosonde Uninhabited Aerial Vehicle (UAV) in the Fourth Convection and Moisture Experiment (CAMEX 4)," *AIAA's 1st Technical Conference and Workshop on Unmanned Aerospace Vehicles*, Portsmouth, VA, May 2002.
- [5] Lin, P.-H., and Lee, C.-S., "Fly into Typhoon Haiyan with UAV Aerosonde," *12th Symposium on Meteorological Observations and Instrumentation*, American Meteorological Society, Long Beach, CA, 2003.
- [6] Lin, P.-H., and Lee, C.-S., "The Eyewall-Penetration Reconnaissance Observation of Typhoon Longwang (2005) with Unmanned Aerial Vehicle, Aerosonde," *Journal of Atmospheric and Oceanic Technology*, Vol. 25, No. 1, 2008, pp. 15–25.

- doi:10.1175/2007JTECHA914.1
- [7] LaValle, S. M., *Planning Algorithms*, Cambridge Univ. Press, New York, 2006.
 - [8] Mitchell, I. M., "The Flexible, Extensible and Efficient Toolbox of Level Set Methods," *Journal of Scientific Computing*, Vol. 35, Nos. 2–3, 2007, pp. 300–329.
 - [9] Philippsen, R., and Siegwart, R., "An Interpolated Dynamic Navigation Function," *Proceedings: IEEE International Conference on Robotics and Automation*, IEEE Publications, Piscataway, NJ, 2005, pp. 3782–3789.
 - [10] Petres, C., Pailhas, Y., Patron, P., Petillot, Y., Evans, J., and Lane, D., "Path Planning for Autonomous Underwater Vehicles," *IEEE Transactions on Robotics and Automation*, Vol. 23, No. 2, 2007, pp. 331–41.
 - [11] Jarvis, R. A., "Collision-Free Trajectory Planning Using Distance Transforms," *Mechanical Engineering, Journal of the Inst. of Engineers*, Vol. ME10, No. 3, 1985, pp. 187–191.
 - [12] Dorst, L., and Trovato, K. I., "Optimal Path Planning by Cost Wave Propagation in Metric Configuration Space," *Proc. SPIE, Mobile Robots III*, edited by W. J. Wolfe, Vol. 1007, International Society for Optical Engineering, Bellingham, WA, Jan. 1988, p. 186.
 - [13] Barraquand, J., Langlois, B., and Latombe, J. C., "Numerical Potential Field Techniques for Robot Path Planning," *IEEE Transactions on Systems, Man and Cybernetics*, Vol. 22, No. 2, April 1992, pp. 224–241.
 - doi:10.1109/21.148426
 - [14] Soullignac, M., Taillibert, P., and Rueher, M., "Adapting the Wave Front Expansion in Presence of Strong Currents," *IEEE International Conference on Robotics and Automation*, IEEE Publications, Piscataway, NJ, 2008, pp. 1352–8.
 - [15] Sethian, J. A., and Vladimirsky, A., "Ordered Upwind Methods for Static Hamilton–Jacobi Equations: Theory and Algorithms," *SIAM Journal on Numerical Analysis*, Vol. 41, No. 1, 2003, pp. 325–363. doi:10.1137/S0036142901392742
 - [16] Straka, J. M., and Anderson, J. R., "The Numerical Simulations of Microburst Producing Thunderstorms: Some Results from Storms Observed During the Cohmex Experiment," *Atmospheric Sciences Rept.*, Vol. 50, No. 10, 1993, pp. 1329–1348. doi:10.1175/1520-0469(1993)050<1329:NSOMPS>2.0.CO;2
 - [17] Gilmore, M. S., and Straka, J. M., "Precipitation Uncertainty Due to Variations in Precipitation Particle Parameters Within a Simple Microphysics Scheme," *Monthly Weather Review*, Vol. 132, No. 11, 2004, pp. 2610–2627. doi:10.1175/MWR2810.1
 - [18] Gilmore, M. S., Straka, J. M., and Rasmussen, E. N., "Precipitation and Evolution in Simulated Deep Convective Clouds Between Liquid-Only and Simple Ice and Liquid Phase Microphysics," *Monthly Weather Review*, Vol. 132, No. 8, 2004, pp. 1897–1916. doi:10.1175/1520-0493(2004)132<1897:PAESIS>2.0.CO;2